

PROPERTIES OF A CENTER/SURROUND RETINEX

PART ONE: SIGNAL PROCESSING DESIGN

Zia-ur Rahman
Science and Technology Corporation
101 Research Drive, Hampton, VA 23666

ABSTRACT

The last version of Edwin Land’s retinex model for human vision’s lightness and color constancy has been implemented. Previous research has established the mathematical foundations of Land’s retinex but has not examined specific design issues and their effects on the properties of the retinex operation. Here we describe the signal processing design of the retinex. We find that the placement of the logarithmic function is important and produces best results when placed after the surround formation. We also find that best rendition is obtained for a “canonical” gain-offset applied after the retinex operation.

1. Introduction

Of the many visual tasks accomplished so gracefully by human vision, one of the most fundamental and approachable for machine vision applications is lightness and color constancy. While a completely satisfactory definition is lacking, lightness and color constancy refer to human perception’s resilience to wide ranging intensity and spectral illumination variations (scene-to-scene and to a large extent within scene). Various theories for this have been proposed and have a common mathematical foundation¹. The last version of Edwin Land’s retinex² has captured our attention because of the ease of implementation and manipulation of key variables, and because it does not have “unnatural” requirements for scene calibration. Likewise, the simplicity of the computation was appealing and initial experiments produced compelling results. This version of the retinex has been the subject of previous digital simulations that were limited because of lengthy computer time involved and was implemented in analog VLSI to achieve real-time computation^{3, 4}. Evidence that this retinex version is an optimal solution to the lightness problem has come from experiments posing Land’s Mondrian target (randomly arranged two-dimensional gray patches) as a problem in linear optimization and a learning problem for back propagated artificial neural networks^{5, 6}.

The utility of a lightness-color constancy algorithm for machine vision is the simultaneous accomplishment of 1) dynamic range compression, 2) color independence from the spectral distribution of the scene illuminant, and 3) color and lightness rendition. Land's center/surround retinex demonstrably achieves the first two although Land emphasized primarily the color constancy properties. Well known difficulties arise though, for color and lightness rendition^{1, 3, 6}. These consist of 1) lightness and color "halo" artifacts that are especially prominent where large uniform regions abut to form a high contrast edge with "graying" in the large uniform zones in an image, and 2) global violations of the gray world assumption (e.g., an all-red scene) which result in a global "graying out" of the image. Clearly the retinex (perhaps like human vision) functions best for highly diverse scenes and poorest for impoverished scenes. This is analogous to systems of simultaneous equations where a unique solution exists if and only if there are enough independent equations.

The general form of the center/surround retinex (Fig. 1) is similar to the difference-of-Gaussian (DOG) function widely used in natural vision science to model both the receptive fields of individual neurons and perceptual processes. The only extensions required are 1) to greatly enlarge and weaken the surround Gaussian (as determined by its space and amplitude constants), and 2) to include a logarithmic function to make subtractive inhibition into a shunting inhibition (i.e., arithmetic division). We have chosen a Gaussian surround form whereas Land opted for a $1/r^2$ function² and Moore et al.³ used a different exponential form. These will be compared in Section 2. Mathematically this takes the form,

$$R_i(x, y) = \log I_i(x, y) - \log [F(x, y) * I_i(x, y)] \quad (1)$$

where $I_i(x, y)$ is the image distribution in the i th color spectral band, "*" denotes the convolution operation, $F(x, y)$ is the surround function, $R_i(x, y)$ is the associated retinex output. Equivalently

$$R_i(x, y) = \log I_i(x, y) - \log \left[\mathcal{F}^{-1} \{ \hat{F}(v, \omega) \cdot \hat{I}_i(v, \omega) \} \right] \quad (2)$$

where $\hat{F}(v, \omega)$ and $\hat{I}_i(v, \omega)$ are the Fourier transforms of $F(x, y)$ and $I_i(x, y)$ and \mathcal{F}^{-1} denotes the inverse Fourier transform.

This operation is performed on each spectral band to produce Land's triplet values specifying color and lightness. It is readily apparent that color constancy (i.e., independence from single source illuminant spectral distribution) is reasonably complete since

$$I_i(x, y) = S_i(x, y) r_i(x, y) \quad (3)$$

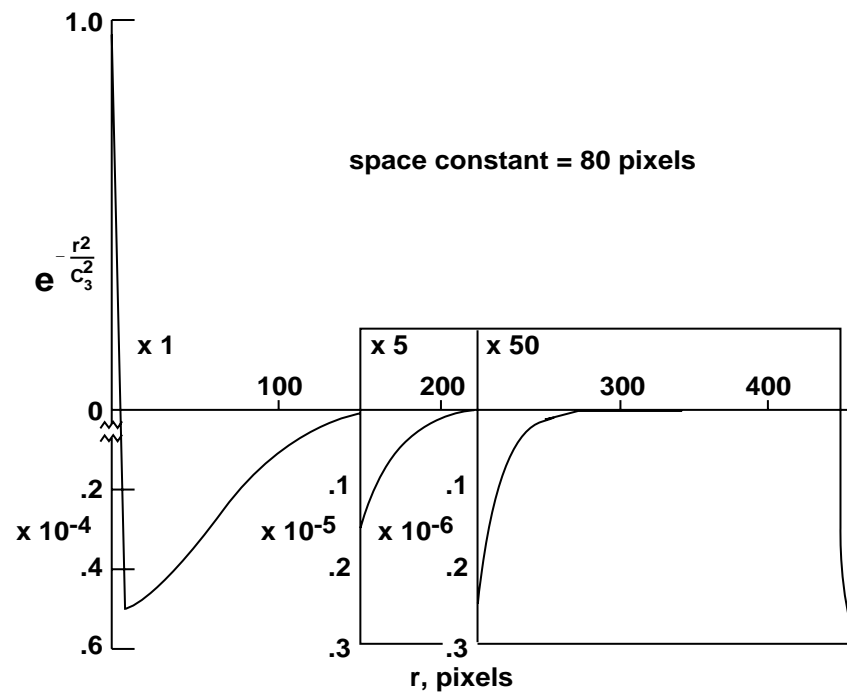
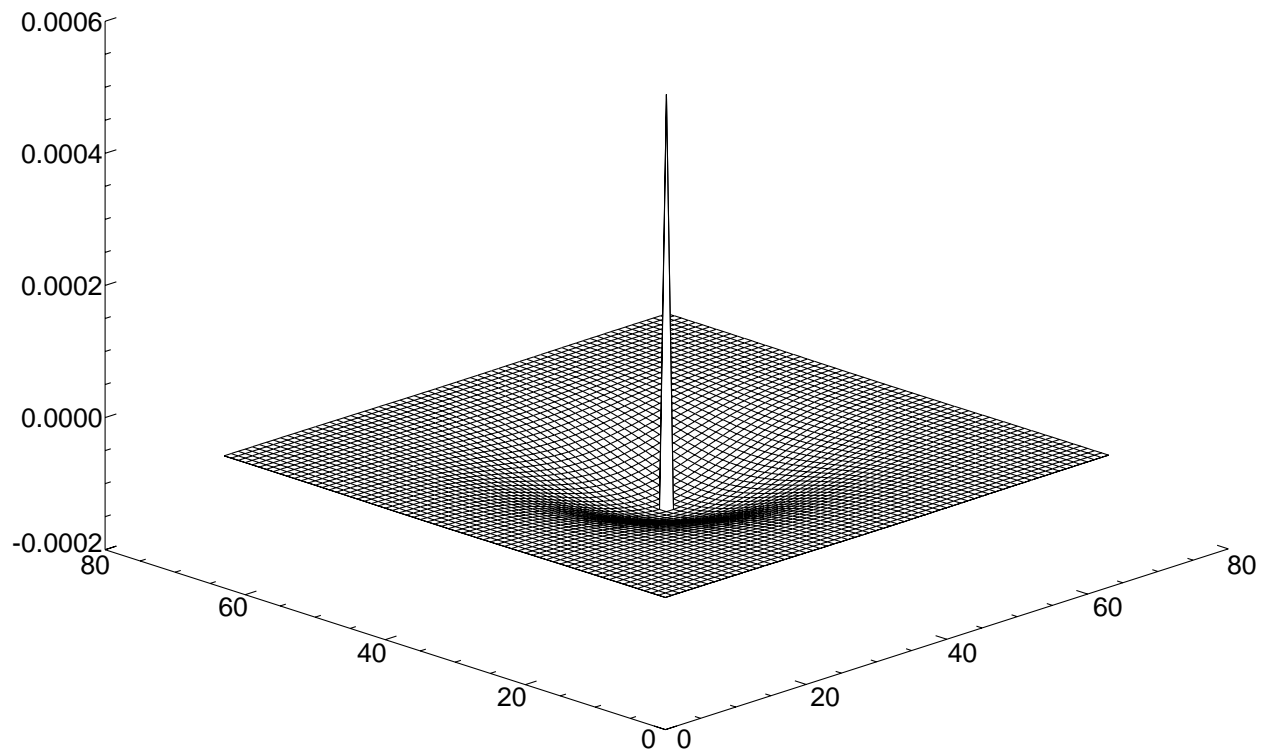


Figure 1: The spatial form of the center/surround retinex operator. a) 3-D representation (distorted to visualize surround). b) Cross-section to illustrate wide weak surround.

where $S_i(x, y)$ is the spatial distribution of the source illumination and $r_i(x, y)$, the distribution of scene reflectances, so that

$$R_i(x, y) = \log \frac{S_i(x, y) r_i(x, y)}{\bar{S}_i(x, y) \bar{r}_i(x, y)} \quad (4)$$

where the bars denote the spatially weighted average value. As long as $S_i(x, y) \approx \bar{S}_i(x, y)$ then

$$R_i(x, y) \approx \log \frac{r_i(x, y)}{\bar{r}_i(x, y)}. \quad (5)$$

The approximate relation is an equality for many cases and, for those cases, where it is not strictly true, the reflectance ratio should dominate illumination variations.

This is demonstrated (Fig. 2) for the extreme cases of blue skylight illumination, direct sunlight only, and tungsten illumination. Actual daylight illumination should fall arbitrarily somewhere between the first two cases. Film and electronic cameras without computational intervention or film selection would produce the top row of images. Dynamic range compression is also readily demonstrated (Fig. 2(right)) with computer simulation. Here the original image data is multiplied by a hyperbolic tangent “shadow.” Again cameras without computation produce the upper result (or with a change of f/stop or exposure would bring out the shadowed detail but at the expense of saturating the non-shadowed image zones). Strikingly, color balance is retained across the wide dynamic range encompassed and the highly nonlinear operation of the retinex. This result is especially provocative to one of us (DJJ) who participated in the colorimetric studies⁷ of the Viking Lander images. The experience with Mars surface color underlined the importance of linear systems methods and the delicacy of color balance which was so easily and dramatically distorted when nonlinearities were employed inappropriately. From this, it would seem that color rendition and nonlinear methods are as immiscible as oil and water. Here, though, is a case where color balance is achieved more automatically with a nonlinear operation. These two examples do, however, point to the difficulty of realizing satisfactory color rendition in contrast to the ease of achieving color constancy and dynamic range compression. Taken together, this discussion indicates the exciting possibilities that motivated us to engage in more extensive investigation.

The need for dynamic range compression and color constancy, especially if both are accomplished simultaneously by a simple real-time algorithm, is well-known to photographers. Discrepancies between the photographer’s perception through the viewfinder and the captured film image can be quite bizarre (Fig. 3) and require constant vigilance to avoid impossible lighting situations, and to carefully select the appropriate film and processing for the illuminant’s spectral distribution. The fundamental limit³ is recognized to be the film or CRT’s narrow dynamic range and static

Figure 2: Demonstration of retinex color constancy and dynamic range compression (prior to optimizing rendition) for a small space constant (space constant = 15 pixels).

spectral response. Print/display dynamic range constraints of 50:1 are, however, compatible with the magnitude of scene reflectance variations. Except for extreme cases (snow or lampblack) reflectance variations are only 20:1⁸ and often much less. Thus even the extremes of reflectance of $\approx 50:1$ are easily spanned by print/display media. Clearly illumination variations are the culprit which human visual perception has overcome by eye-brain computation. Electronic still and video cameras have an intrinsically high dynamic range ($> 2000:1$)⁹ set by the detector array electronics and an even higher dynamic range within the detector array proper, since the limiting factor is usually the preamplifier noise added in transferring image signals off-chip or digitization noise added subsequently. Therefore, at least for electronic cameras, we can conclude that sufficient dynamic range is available to retain the full variations of both illumination and reflectance in arbitrary scenes. So it is certainly reasonable to consider either analog³ implementations of compression/constancy or digital implementation if the initial A/D conversion is done at 10–14 bits, rather than the usual 8 bits.

Recent advances in high speed computing led us to reconsider both extensive digital simulations of the retinex and real time digital implementations for practical use in future electronic camera systems. The hours of computer time previously reported³ are now reduced to minutes and real-time implementations using specialized digital hardware such as digital signal processing (DSP) chips are reasonable. In other words, the full image dynamic range is available from current electronic cameras, real time computation is realizable, and the ultimate bottleneck is only at the first print/display. Obviously, there are image coding aspects to both dynamic range compression and color constancy. Here, we will concentrate on the design of the algorithm to produce combined dynamic range compression/color constancy/color-lightness rendition.

We have seen that the center/surround retinex is both color constant and capable of a high degree of dynamic range compression. It remains then, to specify an implementation that produces satisfactory rendition and examine alternatives to determine if other design options are equally good or better. Because the retinex exchanges illumination variations for scene reflectance context dependency, scene content becomes a major issue especially when it deviates from regionally gray average values—the “gray world” assumption¹. Therefore testing with diverse scenes, including random ones, is important to pinpoint possible limits to the generality of this retinex.

Initial image processing simulations revealed several unresolved implementation issues—1) the placement of the log function, 2) the functional form of the surround, 3) the space constant for the

Figure 3: Examples of serious photographic defects due to spectral and/or spatial illumination variations. a) “Green” kitchen due to fluorescent illumination. b) Sodium vapor illumination. c) Tungsten indoors/daylight outdoors. d) Obscured foreground.

surround, and 4) the treatment of the retinex triplets prior to display. These will now be explored more comprehensively. The results of testing the optimized algorithm on diverse scenes will then be presented with special emphasis on “gray world” violations. Finally, the relationship of the algorithm to neurophysiology will be examined briefly.

2. Signal Processing Issues

A. Placement of Log Function

Previous research^{3, 6} has largely concluded that the logarithm can be taken before or after the formation of the surround. Processing schemes^{3, 6, 10} adhering closely to natural vision science, i.e., an approximate log photoreceptor response, favor placing log response at the photodetection stage prior to any surround formation. Our preliminary testing of this produced rather disappointing results and prompted us to reopen this seemingly decided issue. Initial testing produced encouraging results with much less emphatic artifacts. Mathematically,

$$R_1 = \log I(x, y) - \log[I(x, y) * F(x, y)] \quad (6)$$

and

$$R_2 = \log I(x, y) - \{[\log I(x, y)] * F(x, y)\} \quad (7)$$

are not equivalent. The discrete convolution $[\log I(x, y) * F(x, y)]$ is, in fact, equivalent to a weighted product of $I(x, y)$ whereas the second term in Eq. 6 is a weighted sum. This is closely related to the difference between the arithmetic mean and the geometric mean except that $F(x, y)$ is selected so that

$$\iint F(x, y) dx dy = 1 \quad (8)$$

which does not produce exactly the n th root of n numbers as the geometric mean would. Since the entire purpose of the log operation is to produce a point by point ratio to a large regional mean value, Eq. 6 seems the desired form and our image processing experiments bear out this preference. A typical example is shown in Fig. 4. While the halo artifact for Eq. 7 can be diminished by manipulation of the gain and offset, this results in a significant desaturation of color. In other examples, more severe color distortions occur which likewise cannot be removed by manipulation of the gain-offset. In addition, a shadow simulation indicates much less dynamic range compression for Eq. 7. Therefore, we have selected the Eq. 6 form for our testing and optimization. This form

is also that given in Land’s original presentation² though he is quoted as feeling the two forms were equally useful in practice⁶.

B. Treatment of Retinex Output Prior to Display

During initial experiments we were surprised to find a characteristic form for the histograms of diverse scenes after the retinex operation (Fig. 5). Exceptions were for severe violations of the “gray world” assumption, e.g., an all-red scene. These violations are explored in a subsequent section, so here we will examine a natural image with reasonable scene diversity.

Land’s² proposal of the center/surround retinex does not explicitly address the issue of a final treatment with the possible implication that none is necessary. On the other hand, Moore³ advocates the automatic gain-offset approach whereby the triplet retinex values are adjusted by the absolute maximum and minimum found across all values in all the color bands. Our own empirically-derived approach (Fig. 5) differs from either of these in that a single gain-offset is selected which actually clips both the highest and lowest signal transitions. Little information is lost because the retinex output signals form, to a large degree, a contrast image (being in essence a ratio). Our approach, otherwise, agrees with Moore’s in that a final gain and offset is uniformly applied to all pixels in all three color bands. A comparison of these two approaches is illustrated (Fig. 6) to underline the considerable visual differences encountered. We speculate that the significant deviations from the characteristic histogram which occur for gross violations of the grey world assumption could be used to detect errors. The single gain-offset seems otherwise to be invariant from image to image, so that we have the sense that it may, in fact, be canonical and, therefore, satisfies the original intent of Land to produce a general computation that applies to most images.

C. Conclusions

The specific implementation we have defined from preliminary testing is a center/surround operation with the following characteristics:

1. The logarithm is applied after surround formation by two-dimensional spatial convolution.
2. A “canonical” gain-offset is applied to the retinex output which, in signal terms, clips some of the highest and lowest signal excursions.

Figure 4: Demonstration of improved rendition obtained applying the log response after surround formation.

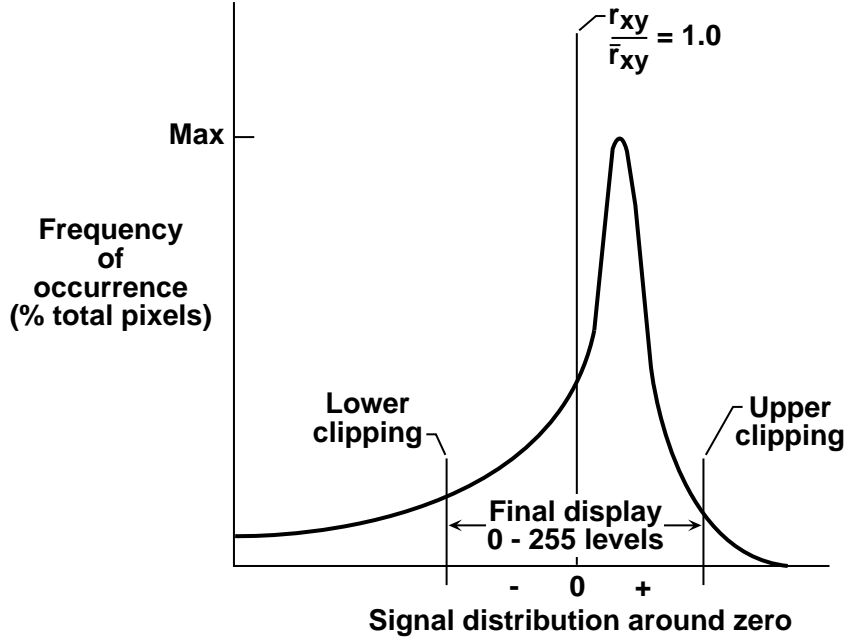


Figure 5: Schematic of a characteristic retinex histogram illustrating the final gain/offset selection applied uniformly to the three color sub-images.

In a companion paper,¹¹ additional components of the design are defined:

3. The form of the surround is Gaussian.
4. The spatial extent of the surround is that for a Gaussian space constant of about 80 pixels (which corresponds to a FWHM spread of 210 pixels).

Our implementation differs from previous ones in that Land proposed an inverse square surround while Moore and Hurlbert concentrated on placement of the log prior to surround formation (or else considered placement as interchangeable). Finally, Moore specified an automatic gain-offset process rather than the canonical one used here. All of these differences were shown to result in significant visual effects on processed images.

Figure 6: Comparison of the visual performance of auto gain/offset versus “canonical” gain/offset. The auto gain/offset is selected on the absolute maximum and minimum values in all three color bands and applied uniformly to all three as a global operation. The “canonical” gain-offset accepts some clipping of extreme high and low values but provides superior rendition with minimal loss of visual information.

REFERENCES

1. A. C. Hurlbert, “Formal connections between lightness algorithms,” *Journal of the Optical Society of America A*, vol. 3, pp. 1684–1693, 1986.
2. E. Land, “An alternative technique for the computation of the designator in the retinex theory of color vision,” *Proc. Nat. Acad. Sci.*, vol. 83, pp. 3078–3080, 1986.
3. A. Moore, J. Allman, and R. M. Goodman, “A real-time neural system for color constancy,” *IEEE Transactions on Neural Networks*, vol. 2, pp. 237–247, March 1991.
4. A. Moore, G. Fox, J. Allman, and R. M. Goodman, “A VLSI neural network for color constancy,” in *Advances in Neural Information Processing 3* (D. S. Touretzky and R. Lippman, eds.), pp. 370–376, San Mateo, CA: Morgan Kaufmann, 1991.
5. A. C. Hurlbert and T. Poggio, “Synthesizing a color algorithm from examples,” *Science*, vol. 239, pp. 482–485, 1988.
6. A. C. Hurlbert, *The Computation of Color*. PhD thesis, Massachusetts Institute of Technology, September 1989.
7. F. O. Huck, D. J. Jobson, S. K. Park, S. D. Wall, R. E. Arvidson, W. R. Patterson, and W. D. Benton, “Spectrophotometric and color estimates of the viking lander sites,” *J. of Geo. Res.*, vol. 28, no. 8, pp. 4401–4411, 1977.
8. D. E. Bowker, R. E. Davis, D. L. Myrick, K. Stacy, and W. L. Jones, “Spectral reflectances of natural targets for use in remote sensing studies,” *NASA References Publication* 1139, June 1985.
9. R. H. Dyck, “Design, fabrication, and performance of ccd imagers,” in *VLSI Electronics Microstructures Science* (N. G. Einspruch, ed.), vol. 3, pp. 65–107, Orlando, FL: Academic Press, 1982.
10. C. Mead, *Analog VLSI and Neural Systems*. Reading, MA: Addison-Wesley, 1989.
11. D. J. Jobson and G. A. Woodell, “Properties of a center/surround Retinex Part Two: Surround design,” *NASA Technical Memorandum* 110188, 1995.